arXiv:1211.2679v3 [stat.AP] 20 Nov 2012

High Dimensional Principal Component Scores and Data Visualization

BY DAN SHEN[1], HAIPENG SHEN[2], HONGTU ZHU[1], AND J. S. MARRON[1,2]

[1]*Department of Biostatistics, University of North Carolina at Chapel Hill, Chapel Hill, North Carolina 27599, U.S.A.*

[2]*Department of Statistics and Operations Research, University of North Carolina at Chapel Hill, Chapel Hill, North Carolina 27599, U.S.A.*

dshen@email.unc.edu *haipeng@email.unc.edu*
htzhu@email.unc.edu *marron@email.unc.edu*

SUMMARY

Principal component analysis is a useful dimension reduction and data visualization method. However, in high dimension, low sample size asymptotic contexts, where the sample size is fixed and the dimension goes to infinity, a paradox has arisen. In particular, despite the useful real data insights commonly obtained from principal component score visualization, these scores are not consistent even when the sample eigenvectors are consistent. This paradox is resolved by asymptotic study of the ratio between the sample and population principal component scores. In particular, it is seen that this proportion converges to a non-degenerate random variable. The realization is the same for each data point, i.e. there is a common random rescaling, which appears for each eigen-direction. This then gives inconsistent axis labels for the standard scores plot, yet the relative positions of the points (typically the main visual content) are consistent. This paradox disappears when the sample size goes to infinity.

Some key words: Principal Component Analysis; High Dimension; Low Sample Size; Spike Model.

1. INTRODUCTION

Visualization of high dimension, low sample size data by principal component analysis has proven to be very useful. A recent example is shown in Figure 1, which studies Next Generation Sequencing for a single gene, in a cancer study, from The Cancer Genome Atlas (TCGA, 2005). The data objects are $n = 180$ curves (each from one biological tissue sample), reflecting the log base 10 read depth, at around $d = 1700$ genome map locations. Relative positions of these biological samples are visualized, using a standard principal components scores plot, in Panel (A) of Figure 1. The plot shows the projection of the data onto the subspace generated by the first two eigenvectors. Note that there is distinct visual impression of three clusters. To investigate this clustering, the clusters have been manually brushed, with three different colors, as shown. To investigate whether these clusters represent important scientific phenomena, the same coloring is applied to the raw data curves in Panel (B). The distinct blocks in Panel (B) represent different exonic regions of the genome, and the jumps in the curves at the boundaries of these blocks indicate splicing events. The red curves are generally very low (recall the log scale) indicating very low levels of expression of this gene, for these samples. The black and blue curves are

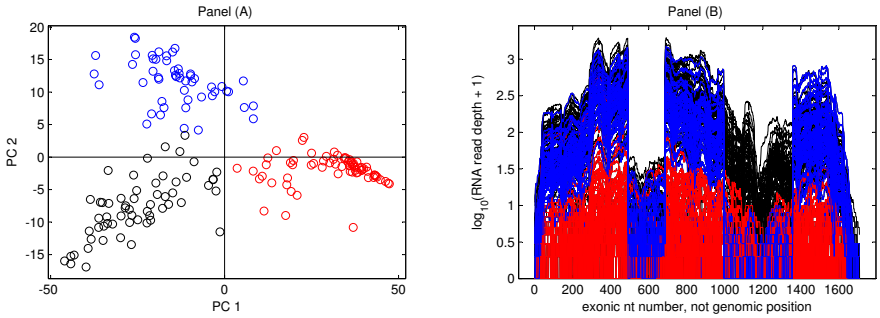

Fig. 1. Principal component analysis of a Next Generation Sequencing cancer data set. The scores plot in Panel (A) suggests three clusters, which are manually brushed with colors. Relevance of these visually discovered clusters is studied in Panel (B), by showing curves with the brushed colors, which reveals a biologically important alternate splicing event.

generally much higher, showing much stronger gene expression. The black and blue curve differ strongly over the fourth exonic region, between 1000 and 1400, where the blue samples show a clear deletion of this exon. Finding such deletions is an important goal in cancer research, as they can form the basis of targeted treatments. Note that this is just one example, where scientifically important structure in data has been discovered by principal component analysis in a high dimension, low sample size context.

We are interested in investigating the mathematical underpinnings of this visual approach to data analysis demonstrated in Figure 1. There are several approaches to this in the literature. Given the nature of genetic data, we prefer to study high dimension, low sample size asymptotics. This approach considers increasing dimension $d \to \infty$ for a fixed sample size n. It has recently been studied in various multivariate analysis contexts, including geometric representation of high dimensional data (Hall et al., 2005), clustering (Ahn et al., 2012), and principal component analysis (Ahn et al., 2007; Jung & Marron, 2009; Jung et al., 2012; Yata & Aoshima, 2012; Shen et al., 2012a). However, these asymptotic analyses of principal component analysis all focused on studying the angles between the sample eigenvectors and the corresponding population eigenvectors. For example, under some mild conditions, Jung & Marron (2009) showed that such angles go to 0, which is defined as the *consistency* of the sample eigenvectors.

In this paper, we take a deeper look by studying principal component scores, shown as the circles in Panel (A) of Figure 1. Our analysis surprisingly reveals an apparent paradox under the high dimension low sample size setting: principal component scores are *inconsistent* with the corresponding population scores, even when the sample eigenvectors are consistent. Furthermore, for a fixed n and a particular principal component, as d goes to infinity, the proportion between the sample scores and the corresponding population scores converges to a random variable, whose realization is the same for each observation. The findings suggest that, although the principal component scores can not be consistently estimated, the scores scatter plots, such as Panel (A) of Figure 1, can still be used to explore interesting features of high dimension, low sample size data, because the relative positions of the points are consistent due to the common scaling. Finally, this phenomenon disappears when the sample size tends to infinity. In particular, both the sample eigenvectors and the sample principal component scores are then consistent.

2. Notations and Assumptions

Assume that $X_1, \ldots, X_n$ are a random sample from the d-dimensional normal distribution $N(\xi, \Sigma)$, and the population covariance matrix Σ has the following eigen-decomposition:

$$\Sigma = U \Lambda U^T,$$

where Λ is the diagonal matrix of the population eigenvalues $\lambda_1 \geq \ldots \geq \lambda_d$, and U is the corresponding eigenvector matrix such that $U = [u_1, \ldots, u_d]$. Denote the jth normalized population principal component score vector as

$$S_j = (S_{1,j}, \cdots, S_{n,j})^T \equiv \lambda_j^{-\frac{1}{2}} (u_j^T X_1, \cdots, u_j^T X_n)^T, \quad j = 1, \cdots, d. \tag{1}$$

Let $\overline{X}$ be the sample mean. As discussed in Paul & Johnstone (2007),

$$\sum_{i=1}^{n} (X_i - \overline{X})(X_i - \overline{X})^T \quad \text{has the same distribution as} \quad \sum_{i=1}^{n-1} Y_i Y_i^T,$$

where Y_i are independent and identically distributed random variables from $N(0, \Sigma)$. It follows that the sample covariance matrix is location invariant. Without loss of generality, we assume that $X_1, \ldots, X_n$ are a random sample from the d-dimensional normal distribution $N(0, \Sigma)$.

Denote the data matrix as $X = [X_1, \ldots, X_n]$, and the sample covariance matrix as $\hat{\Sigma} = n^{-1} X X^T$, which has the following eigen-decomposition,

$$\hat{\Sigma} = \hat{U} \hat{\Lambda} \hat{U}^T, \tag{2}$$

where $\hat{\Lambda} = \text{diag}(\hat{\lambda}_1, \cdots, \hat{\lambda}_d)$ is the diagonal sample eigenvalue matrix, and $\hat{U} = [\hat{u}_1, \ldots, \hat{u}_d]$ is the corresponding sample eigenvector matrix. In addition, the matrix $X/\sqrt{n}$ has the following singular value decomposition: $X/\sqrt{n} = \sum_{j=1}^{d} \hat{\lambda}_j^{\frac{1}{2}} \hat{u}_j \hat{v}_j$, where $\hat{v}_j = (\hat{v}_{1,j}, \cdots, \hat{v}_{n,j})$, $j = 1, \cdots, d$. Then the jth normalized sample principal component vector is

$$\hat{S}_j = (\hat{S}_{1,j}, \cdots, \hat{S}_{n,j})^T = (\hat{v}_{1,j}, \cdots, \hat{v}_{n,j})^T, \quad j = 1, \cdots, d. \tag{3}$$

Panel (A) of Figure 1 shows a scatter plot of the $\hat{S}_{i,1}$ versus $\hat{S}_{i,2}$, $i = 1, \cdots, n$.

3. Asymptotic Properties of Principal Component Scores

The asymptotic properties of principle component scores in high dimension, low sample size contexts are studied in Section 3·1, and as the sample size grows in Section 3·2.

3·1. *High Dimension, Low Sample Size Analysis*

In this subsection, we consider the high dimension, low sample size settings, where the sample size n is fixed and the dimension d goes to infinity. We consider multiple spike models (Jung & Marron, 2009) under which, as $d \to \infty$,

$$\lambda_1 \gg \cdots \gg \lambda_m \gg \lambda_{m+1} \sim \cdots \sim \lambda_d \sim 1, \tag{4}$$

where $\lambda_i \gg \lambda_j$ means that $\lim_{d\to\infty} \lambda_j/\lambda_i = 0$, and $\lambda_i \sim \lambda_j$ means that $c_1 \leq \underline{\lim}_{d\to\infty} \lambda_i/\lambda_j \leq \overline{\lim}_{d\to\infty} \lambda_i/\lambda_j \leq c_2$ for constants $c_1 \leq c_2$.

Under the above spike models, Jung & Marron (2009) showed that when n is fixed, if $d/\lambda_m \to 0$, the angle between each of the first m sample eigenvectors $\hat{u}_j$ and its corresponding population eigenvector u_j goes to 0 with probability 1, which is defined as the *consistency* of the sample eigenvector.

However, under the same assumptions, an anonymous reviewer identified a paradoxical phenomenon in that the sample principal component scores are not consistent. In addition, our analysis suggests that, for a particular principal component, the proportion between the sample principal component scores and the corresponding population scores converges to a random variable, the realization of which remains the same for all data points. These results are summarized in the following Theorem 1. The findings suggest that it remains valid to use score scatter plots as a graphical tool to identify interesting features in high dimension low sample size data.

THEOREM 1. *Under Assumption* (4)*, and for the fixed* n*, as* $d \to \infty$*, if* $d/\lambda_m \to 0$*, then the proportion between the sample and population principal component scores satisfies*

$$\left| \frac{\hat{S}_{i,j}}{S_{i,j}} \right| \xrightarrow{p} R_j, \quad i = 1, \cdots n, \ j = 1, \cdots, m, \tag{5}$$

where $\xrightarrow{p}$ *stands for convergence in probability, and* R_j *has the same distribution as* $\sqrt{n/\chi^2_n}$ *with* χ^2_n *being the Chi-square distribution with* n *degrees of freedom.*

Remark 1. Under the assumptions in Theorem 1, Jung & Marron (2009) and Shen et al. (2012b) have shown that the angle between the sample eigenvector $\hat{u}_j$ and the corresponding population eigenvector u_j, for $j = 1, \cdots, m$, converges to 0 with probability 1, which suggests that the sample eigenvectors are consistent, although the principal component scores are inconsistent under the same assumptions.

Remark 2. It follows from (5) that the ratio R_j only depends on j (the index of the principal components), but not i (the index of the data points). This particular scaling suggests that the scores scatter plot, such as Panel (A) of Figure 1, has incorrectly labeled axes (by the common factor R_j for the corresponding axis), and yet asymptotically correct relative positions of the points; hence the scatter plot still enables meaningful identification of useful scientific features as demonstrated in Panel (B).

3·2. *Growing Sample Size Analysis*

In this subsection, we consider growing sample size contexts, where $n \to \infty$, and then study the asymptotic properties of the principal component scores. This includes a wide range of settings, including classical asymptotics, where dimension d is fixed, random matrix asymptotics where $d \sim n$ and more, see Shen et al. (2012b) for an overview. Unlike the low sample size setting, the apparent inconsistency paradox now disappears. This means that both the sample eigenvectors and the sample principal component scores can be consistent.

We consider the following multiple spike models, as $n \to \infty$,

$$\lambda_1 \succ \cdots \succ \lambda_m \gg \lambda_{m+1} \sim \cdots \sim \lambda_d \sim 1. \tag{6}$$

Here $\lambda_i \succ \lambda_j$ means that $\overline{\lim}_{n\to\infty} \lambda_j/\lambda_i < 1$. Compared with the multiple spike models (4), the multiple spike assumption (6) is weaker because we have more sample information ($n \to \infty$).

Theorem 2 suggests that as $n \to \infty$, the proportion between the sample scores and the corresponding population scores tends to 1. This connects with the above results, from the fact that the ratio R_j in (5) has the same distribution as an asymptotic $\sqrt{n/\chi^2_n}$ distribution which converges almost surely to 1 as $n \to \infty$. Thus, it is not surprising that the apparent inconsistency disappears as the sample size grows.

THEOREM 2. *Under Assumption* (6)*, and as $n \to \infty$, if $d/\lambda_m \to 0$, then the proportion between the sample and population principal component scores satisfies*

$$\left| \frac{\hat{S}_{i,j}}{S_{i,j}} \right| \xrightarrow{a.s} 1, \quad i = 1, \cdots n, \ j = 1, \cdots, m, \tag{7}$$

where $\xrightarrow{a.s}$ stands for almost sure convergence.

Remark 1. Under the current context, the consistency of the sample principal component scores fits as expected, with the fact that the sample eigenvectors are consistent under the assumptions of Theorem 2. In particular, Shen et al. (2012b) have shown that, under the same assumptions, the angle between the sample eigenvector $\hat{u}_j$ and the corresponding population eigenvector u_j for $j = 1, \cdots, m$ converges almost surely to 0.

APPENDIX

In this section, we provide the technical details of the proofs for Theorems 1 and 2. First, we present two lemmas from Shen et al. (2012b), that will be used to prove Theorems 1 and 2.

LEMMA 1. *Under the assumptions in Theorem* 1 *and as $d \to \infty$, the sample eigenvalues satisfy*

$$\frac{\hat{\lambda}_j}{\lambda_j} \xrightarrow{p} \frac{\chi^2_n}{n}, \quad j = 1, \cdots, m, \tag{A1}$$

and the sample eigenvectors satisfy

$$\begin{cases} \lambda_k^{\frac{1}{2}} \lambda_j^{-\frac{1}{2}} |\hat{u}_j^T u_k| \xrightarrow{p} 1, \quad \textit{for } 1 \le k = j \le m, \quad \textit{or} \quad 0 \ \textit{for } 1 \le k \ne j \le m, \\ \sum_{k=m+1}^{d} (\hat{u}_j^T u_k)^2 \xrightarrow{p} 0, \quad \textit{for } 1 \le j \le m. \end{cases} \tag{A2}$$

LEMMA 2. *Under the assumptions in Theorem* 2 *and as $n \to \infty$, the sample eigenvalues satisfy*

$$\frac{\hat{\lambda}_j}{\lambda_j} \xrightarrow{a.s} 1, \quad j = 1, \cdots, m, \tag{A3}$$

and the sample eigenvectors satisfy

$$\begin{cases} \lambda_k^{\frac{1}{2}} \lambda_j^{-\frac{1}{2}} |\hat{u}_j^T u_k| \xrightarrow{a.s} 1, \quad \textit{for } 1 \le k = j \le m, \quad \textit{or} \quad 0 \ \textit{for } 1 \le k \ne j \le m, \\ \sum_{k=m+1}^{d} (\hat{u}_j^T u_k)^2 \xrightarrow{a.s} 0, \quad \textit{for } 1 \le j \le m. \end{cases} \tag{A4}$$

Note that X_i has the following decomposition

$$X_i = \sum_{j=1}^{d} \lambda_j^{\frac{1}{2}} u_j z_{i,j}, \tag{A5}$$

where the $z_{i,j}$'s are independent and identically standard normally distributed for $i = 1, \cdots, n$, $j = 1, \cdots, d$. It follows from (1) and (A5) that the jth population principal component scores are

$$S_j = (S_{1,j}, \cdots, S_{n,j})^T = (z_{1,j}, \cdots, z_{n,j})^T. \tag{A6}$$

From (3), the jth sample principal component scores are

$$\hat{S}_j = (\hat{S}_{1,j}, \cdots, \hat{S}_{n,j}) = \hat{\lambda}_j^{-\frac{1}{2}} (\hat{u}_j^T X_1, \cdots, \hat{u}_j^T X_n). \tag{A7}$$

From (A5), (A6) and (A7), we have that the proportion between the sample principal component scores and the corresponding population scores are

$$\frac{\hat{S}_{i,j}}{S_{i,j}} = \frac{\lambda_j^{\frac{1}{2}}}{\hat{\lambda}_j^{\frac{1}{2}}} \hat{u}_j^T u_j + \sum_{1 \le k \le m, k \ne j}^{m} \frac{\lambda_k^{\frac{1}{2}} z_{i,k}}{\hat{\lambda}_j^{\frac{1}{2}} z_{i,j}} \hat{u}_j^T u_k + \sum_{k=m+1}^{d} \frac{\lambda_k^{\frac{1}{2}} z_{i,k}}{\hat{\lambda}_j^{\frac{1}{2}} z_{i,j}} \hat{u}_j^T u_k. \tag{A8}$$

Proof of Theorem 1. It follows from Lemma 1 that as $d \to \infty$

$$|\frac{\lambda_j^{\frac{1}{2}}}{\hat{\lambda}_j^{\frac{1}{2}}} \hat{u}_j^T u_j| \xrightarrow{p} R_j, \qquad \sum_{1 \le k \le m, k \ne j}^{m} \frac{\lambda_k^{\frac{1}{2}} z_{i,k}}{\hat{\lambda}_j^{\frac{1}{2}} z_{i,j}} \hat{u}_j^T u_k \xrightarrow{p} 0, \tag{A9}$$

where R_j has the same distribution as $\sqrt{n/\chi_n^2}$. Without loss of generality, we assume that $\lambda_k = 1$ for $k = m+1, \cdots, d$. Then it follows from Cauchy-Schwarz inequality that

$$\left\{ \sum_{k=m+1}^{d} \frac{\lambda_k^{\frac{1}{2}} z_{i,k}}{\hat{\lambda}_j^{\frac{1}{2}} z_{i,j}} \hat{u}_j^T u_k \right\}^2 \le \frac{d-m}{\hat{\lambda}_j z_{i,j}^2} \left\{ \frac{1}{d-m} \sum_{k=m+1}^{d} z_{i,k}^2 \right\} \left\{ \sum_{k=m+1}^{d} (\hat{u}_j^T u_k)^2 \right\}^2. \tag{A10}$$

From Lemma 1 and (A10), the last item in the right-hand-side of Equation (A8) converges to 0 with probability 1. Combining the above with (A9), we obtain (5). In addition, it follows from (A2) that the angle between $\hat{u}_j$ and u_j converges to 0 with probability 1, which concludes the proof of Theorem 1.

Proof of Theorem 2. The proof of Theorem 2 is similar. To avoid overlap, details are not given here. The critical difference in the proof of Theorem 2 is the use of Lemma 2, i.e. (A1) should be replaced by (A3).

REFERENCES

AHN, J., LEE, M. & YOON, Y. (2012). Clustering high dimension, low sample size data using the maximal data piling distance. *Statistica Sinica* **22**, 443–464.

AHN, J., MARRON, J., MULLER, K. & CHI, Y. (2007). The high-dimension, low-sample-size geometric representation holds under mild conditions. *Biometrika* **94**, 760–766.

HALL, P., MARRON, J. & NEEMAN, A. (2005). Geometric representation of high dimension, low sample size data. *Journal of the Royal Statistical Society: Series B* **67**, 427–444.

JUNG, S. & MARRON, J. (2009). PCA consistency in high dimension. low sample size context. *Annals of Statistics* **37**, 4104–4130.

JUNG, S., SEN, A. & MARRON, J. (2012). Boundary behavior in high dimension, low sample size asymptotics of pca. *Journal of Multivariate Analysis* **109**, 190–203.

PAUL, D. & JOHNSTONE, I. (2007). Augmented sparse principal component analysis for high-dimensional data. *Technical Report, UC Davis* .

SHEN, D., SHEN, H. & MARRON, J. (2012a). Consistency of sparse PCA in high dimension and low sample size contexts. *Journal of Multivariate Analysis, forthcoming* .

SHEN, D., SHEN, H. & MARRON, J. (2012b). A general framework for consistency of principal component analysis: Supplement materials. *Technical Report, UNC-Chapel Hill* .

TCGA (2005). *The Cancer Genome Atlas homepage: "http://cancergenome.nih.gov/"*.

YATA, K. & AOSHIMA, M. (2012). Effective PCA for high-dimension, low-sample-size data with noise reduction via geometric representations. *Journal of Multivariate Analysis* **105**, 193–215.